CLAIMS

What is claimed is:

1. A snapshot tree structure, comprising:

a first branch, comprising:

a base volume storing a current user data;

a first read-only snapshot descending from the base volume, the first read-only snapshot being created at a first time, the first read-only snapshot storing a first data of the base volume at the first time before the first data is modified in the base volume; and

a second read-only snapshot descending from the first snapshot, the second read-only snapshot being created at a second time earlier than the first time, the second read-only snapshot storing a second data of the base volume at the second time before the second data is modified in the base volume.

2. The snapshot tree structure of claim 1, further comprising:

a second branch, comprising a first read-write snapshot descending from one of the first and the second read-only snapshots.

3. The snapshot tree structure of claim 2, wherein the second branch further comprises a third read-only snapshot descending from the first read-write snapshot, the third read-only snapshot being created at a third time, the third read-only snapshot storing a third data of the first read-write snapshot at the third time before the third data is modified in the first read-write snapshot.

4. The snapshot tree structure of claim 3, further comprising:

third branch, comprising a second read-write snapshot descending from the third read-only snapshot.

5. The snapshot tree structure of claim 4, wherein the third branch further comprises a fourth read-only snapshot descending from the second read-write snapshot, the fourth read-only snapshot being created at a fourth time, the fourth read-only snapshot storing a fourth data of the second read-write snapshot at the fourth time before the fourth data is modified in the read read-write snapshot.

6. A method for generating a snapshot tree structure, comprising:

creating a first branch, comprising:

creating a base volume storing a current user data;

creating a first read-only snapshot descending from the base volume, the first read-only snapshot being created at a first time;

storing in the first read-only snapshot a first data of the base volume at the first time before the first data is modified in the base volume;

creating a second read-only snapshot descending from the base volume, the second read-only snapshot being created at a second time later than the first time;

storing in the second read-only snapshot a second data of the base volume at the second time before the second data is modified in the base volume; and

inserting the second read-only snapshot between the base volume and the first read-only snapshot, wherein the first read-only snapshot now descends from the second read-only snapshot.

7. The method of claim 6, further comprising:

creating a second branch, comprising creating a first read-write snapshot descending from one of the first and the second read-only snapshots.

8. The method of claim 7, wherein said creating a second branch further comprises creating a third read-only snapshot descending from the first read-write snapshot, the third read-only snapshot being created at a third time, the third read-only snapshot storing a third

data of the first read-write snapshot at the third time before the third data is modified in the first read-write snapshot.

9. The method of claim 8, further comprising:

creating a third branch, comprising creating a second read-write snapshot descending from the third read-only snapshot.

10. The method of claim 9, wherein said creating a third branch further comprises creating a fourth read-only snapshot descending from the second read-write snapshot, the fourth read-only snapshot being created at a fourth time, the fourth read-only snapshot storing a fourth data of the second read-write snapshot at the fourth time before the fourth data is modified in the read read-write snapshot.

11. A method for reading a value of a data block from a snapshot tree structure having a base volume, a first snapshot descending from the base volume, and a second snapshot descending from the first snapshot, the method comprising:

searching for the data block in the second snapshot;

if the data block is not found in the second snapshot:

following a link in the second snapshot to the first snapshot;

searching for the data block in the first snapshot.

12. The method of claim 11, wherein the first and the second snapshots are read-only snapshots.

13. The method of claim 11, wherein the first snapshot is a read-only snapshot and the second snapshot is a read-write snapshot.

14. The method of claim 11, further comprising:

if the data block is not found in the first snapshot:

following a link in the first snapshot to the base volume;

reading the data block from the base volume.